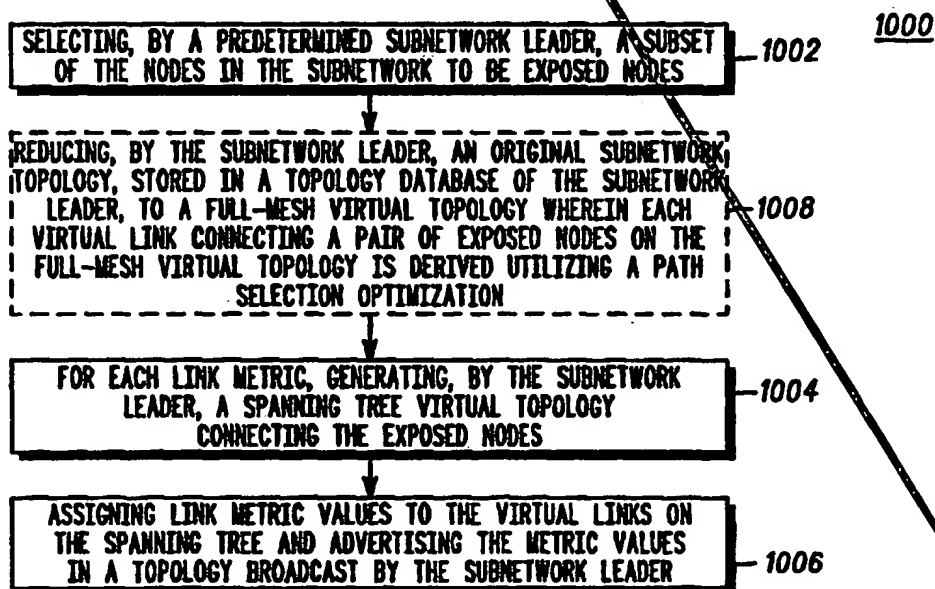




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04L 12/44		A1	(11) International Publication Number: WO 95/31058
			(43) International Publication Date: 16 November 1995 (16.11.95)
(21) International Application Number: PCT/US95/01789		(81) Designated States: AU, CA, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 9 February 1995 (09.02.95)		<p>Published With international search report.</p>	
(30) Priority Data: 08/239,178 6 May 1994 (06.05.94) US			
(71) Applicant: MOTOROLA INC. [US/US]; 1303 East Algonquin Road, Schaumburg, IL 60196 (US).			
(72) Inventor: LEE, Whay, Chiou; Westgate Apartments #1401, Cambridge, MA 02139 (US).			
(74) Agents: STOCKLEY, Darleen, J. et al.; Motorola Inc., Intellectual Property Dept., 1303 East Algonquin Road, Schaumburg, IL 60196 (US).			

(54) Title: METHOD FOR EFFICIENT AGGREGATION OF LINK METRICS



(57) Abstract

The method of the present invention provides efficient aggregation of link metrics for a subnetwork (102, 202, 302) in a communication network (100, 200, 300) having a plurality of interconnected subnetworks (102, 202, 302) that includes nodes (A, B, C, D) and links. Link metric values are summarized on a full-mesh topology (202) and encoded by a spanning tree topology (1004). The method encodes the link metric information contained in full-mesh topology with minimum compromise in the capability to recover it.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

METHOD FOR EFFICIENT AGGREGATION OF LINK METRICS

5

Field of the Invention

This invention relates generally to link state parameters in a communication network, and more particularly to efficient aggregation of link states for a subnetwork in a communication network.

Background

15 In a communication network consisting of nodes representing switching systems and links representing transmission facilities between pairs of nodes, the network topology, i.e., the collection of all link state information, is maintained by each node that is responsible for path computation for establishing communication between nodes. Routing protocols, which are used to select paths, may be broadly divided into two classes: link state protocols and distance vector protocols. Link state protocols require each node to maintain the states of all relevant links. Distance vector protocols require each node to maintain the state of the shortest path from itself to each relevant destination.

Link metrics are quantitative link state parameters that are associated with the quality of information transfer and are used in path computation. There are two kinds of link metrics: non-additive and additive link metrics. A non-additive link metric generically takes the form of "bandwidth". An additive link metric generically takes the form of "delay". The generic path selection

problem is to determine a path according to a given routing objective, such that the bandwidth of each link on the selected path is greater than a given bandwidth threshold, and that the delay along the selected path is smaller than a given delay threshold. There are two common routing objectives in the prior art: maximum bandwidth and minimum delay.

Path metrics are derived from link metrics. For a non-additive link metric such as bandwidth, the metric value associated with a path is the minimum metric value among all the links along the path. For an additive link metric such as delay, the metric value associated with a path is the sum of the metric values of the links along the path.

The topology of the network is kept in a topology database attached to each node that is responsible for maintaining and using such information. Due to changes in the network, the topology must be updated from time to time via a topology broadcast mechanism. A topology broadcast is an event executed by a node such that a message containing link state information is advertised or distributed to all other nodes in the network.

There are two general approaches in the prior art to address topology update complexity. In one approach, the prior art reduces complexity by limiting the handling of the topology to a subset of the nodes in the network. In another approach, the prior art aggregates the link states in each subnetwork so that nodes outside the subnetwork need only maintain partial information on its topology. Typically, a representative node is elected from the nodes in the subnetwork to be responsible for aggregating and

advertising the subnetwork link metrics to nodes outside the subnetwork.

The prior art that makes no compromise in the amount of link state information divides all nodes into two types: (i) master nodes with extended memory and computing capabilities, and (ii) subordinate nodes with limited memory and computing capabilities. A topology database is attached to each master node and is used by the node for path computation. The subordinate nodes rely on the master nodes for topology maintenance and path computation.

In an extreme approach for link metric aggregation, the entire subnetwork is represented by a pseudonode. Such an approach can be traced back to the IS-IS, i.e., Intermediate System-to-Intermediate System, routing protocol. In one version of this approach, all internal link metrics of the subnetwork are reduced to a node metric for the pseudonode by taking worst cases. For an additive link metric, a popular assignment of the corresponding node metric is the "diameter" or the length of the longest shortest path between any pair of nodes in the subnetwork. For a non-additive link metric, a popular assignment of the corresponding node metric is the bandwidth of the smallest maximum-bandwidth path between any pair of nodes in the subnetwork. Another version of the approach is to let the pseudonode to represent the "middle" of the subnetwork, and assign an appropriate metric, by taking averages, to each of the links connecting the pseudonode and the exposed nodes. A variation of the above versions makes use of a diameter variance metric in addition to the diameter metric to reflect deviation from a symmetric topology.

In the OSPF, i.e., the Open Shortest Path First, routing protocol, summary link states are used for inter-subnetwork routing. The summary links are essentially
5 virtual links to external nodes. From the point of view of a subnetwork, there are internal links that connect the internal nodes together, and there are summary links that emanate from the subnetwork to external nodes. OSPF is designed to be a link state routing protocol.

10

In the prior art, a network is clustered into a number of levels, and each node has knowledge of only the shortest paths within its own lowest level cluster and a single shortest path to each supercluster. Path selection is
15 based on a distance-vector method, where a routing table indicating the best next hop for each intended destination is maintained by each node.

Also in the prior art that aggregates link states, the
20 original topology of the subnetwork is reduced to a smaller topology consisting of a subset of the nodes, namely exposed nodes, in the subnetwork. This method advertises a full-mesh virtual topology whose nodes are the exposed nodes and whose links represent paths in the original
25 topology. The paths are determined with respect to a given routing objective. The links in the virtual topology are known as virtual links. This approach is well suited to a hierarchical routing architecture.

30 In other prior art, a different problem is solved. A heuristic that starts by computing a full-mesh shortest-path network consisting of a subset of the nodes in a given original network and then reducing the full-mesh network to a minimum spanning tree is used for routing to multiple

destinations in a computer network, where the object is to minimize the total link weight on the multicast tree that connects the destination nodes together. There are two common spanning tree algorithms known in the prior art:
5 maximum spanning tree algorithm and minimum spanning tree algorithm.

Although the approaches that advertise the topology in terms of a pseudonode, i.e., variations of the approach in
10 IS-IS, offer the greatest reduction of advertised information, these approaches typically do not supply enough information for efficient routing. Moreover, the version that handles worst cases cannot reflect the link state information associated with an asymmetric
15 subnetwork topology. The version that takes averages cannot offer any performance guarantee. The version that uses a diameter variance metric is still under investigation. When a subset of the nodes in the subnetwork are exposed, all versions do not apply.

20

The approach that advertises the full-mesh representation of the subnetwork requires advertisement of order M^2 pieces of link state information when the number of exposed nodes is M . For a subnetwork whose
25 nodes are not very densely connected, the amount of link state information to be advertised may well exceed that contained in the original topology of the subnetwork when M is sufficiently large. Such explosion of link state information contradicts the purpose of link metric
30 aggregation, i.e., efficiency in routing.

The summary link approach in OSPF does not capture enough link state information for efficient routing that is subject to quality of service constraints. Specifically, it

does not offer the flexibility to pick and choose link resources that best meet the quality of service requirements for a connection.

5 Efficient aggregation of link metrics should provide more efficient routing through the communication network. Hence, there is a need for a method for efficient aggregation of link metrics such that the number of link metrics to be advertised is minimized without appreciably
10 compromising the information contained in the link metrics.

Brief Descriptions of the Drawings

15

FIG. 1 is a representation of a topology of a subnetwork for utilizing the method of the present invention.

20

FIG. 2 is a representation of a full-mesh topology of a subnetwork for utilizing the method of the present invention.

25

FIG. 3 is a representation of a spanning tree topology of a subnetwork for utilizing the method of the present invention.

30

FIG. 4 is a representation of exemplary bandwidth and delay metrics in a subnetwork for utilizing the method of the present invention.

FIG. 5 is a representation of exemplary bandwidth and delay metrics for a full-mesh topology derived from

maximum bandwidth paths utilizing the method of the present invention.

FIG. 6 is a representation of exemplary metrics for a maximum bandwidth spanning tree in accordance with the method of the present invention.

FIG. 7 is a representation of exemplary metrics for a maximum delay spanning tree in accordance with the method of the present invention.

FIG. 8 is a representation of exemplary metrics advertised for bandwidth and delay in accordance with the method of the present invention.

FIG. 9 is a representation of exemplary estimates for a full-mesh topology in accordance with the method of the present invention.

FIG. 10 is a flow chart of one embodiment of steps in accordance with the method of the present invention, which provides efficient aggregation of link metrics for a subnetwork in a communication network having a plurality of interconnected subnetworks that consist of nodes and links.

FIG. 11 is a flow chart of one embodiment of steps for aggregation of link metrics for bandwidth maximization utilizing the Maximum Bandwidth Link Metric Aggregation Heuristic in accordance with the method of the present invention.

FIG. 12 is a flow chart of one embodiment of steps for aggregation of link metrics for delay minimization

utilizing the Minimum Delay Link Metric Aggregation Heuristic in accordance with the method of the present invention.

5 FIG. 13 is a flow chart of an embodiment of steps in accordance with the method of the present invention.

FIG. 14 is a flow chart of one embodiment of, in the step of assigning link metric values to the virtual links in
10 accordance with the present invention, selecting a virtual link, on the full-mesh virtual topology, whose metric value is to be estimated from the corresponding spanning tree.

15 **Detailed Description of a Preferred Embodiment**

In a large internetworking environment consisting of many subnetworks, aggregation of link metrics would be advantageous, in particular for two reasons. First, the
20 amount of link state information must be compressed to reduce the topology update complexity which is proportional to the number of links. Second, the actual topology of a subnetwork may have to be hidden for security reasons. The method of the present invention is
25 applicable to link state routing protocols, and provides an efficient solution to the problem of aggregating link states, specifically link metrics, within a subnetwork where a predetermined subset of the nodes are exposed, i.e., to be retained in the process of aggregation. This
30 method requires significantly less communication bandwidth for topology updates than the conventional method where every node in the network is involved in maintaining and updating the entire topology of the network. The method of the present invention may also be

applied recursively to a network with hierarchical structures.

5 The purpose of link metric aggregation is to represent the original topology in the subnetwork by a virtual topology that consists of only the exposed nodes and the smallest number of virtual links that can adequately capture the internal link metrics of the subnetwork for the purpose of path selection based on a
10 given routing objective. There may be one such virtual topology for each link metric. The virtual topologies do not have to be the same since only the link metric information derived from them is of importance.

15 For example, consider aggregating link metrics for a subnetwork with N connected nodes, M of which are exposed, or retained in the aggregation, where $M < N$, where M and N are positive integers. FIG. 1, numeral 100, shows an exemplary subnetwork (102) whose exposed nodes (A, B, C, D) are shaded. For M exposed nodes, the largest number of virtual links that connect them together, since no
20 parallel links are allowed, is $M(M-1)/2$, i.e., order M^2 , and such a fully connected topology, FIG. 2, numeral 200, is known as a full-mesh topology (fully interconnected A-B-C-D; 202). The smallest number of virtual links that
25 connect all the exposed nodes together is $(M-1)$, and any connected virtual topology with $(M-1)$ virtual links belongs to the family of topologies known as "trees". Specifically, since all M exposed nodes are on the tree in this case, the tree, shown in FIG. 3, numeral 300, is known as a "spanning
30 tree". A tree topology consists of no loops. Also, there is a unique path connecting each pair of nodes on the spanning tree.

Where topology update complexity is not an issue, one would clearly want to advertise all the link metric information contained in the full-mesh virtual topology, thereby obtaining the best, i.e., most complete, network topology without permitting parallel links. The representation of the original subnetwork topology by a full-mesh topology connecting a predetermined set of exposed nodes is referred to as a complete summarization of the subnetwork topology. Where selected, the information contained in the complete summarization is compressed by means of encoding. The method of the present invention encodes the link metric information contained in the full-mesh topology, together with minimum compromise in the capability to recover it. The order M^2 pieces of link metric information are compressed to order M pieces. Given this encoded link metric information, one can obtain $(M-1)$ pieces of the link metric information with complete accuracy and derive reasonable estimates of the remaining pieces in the form of lower bounds for bandwidth metrics and upper bounds for delay metrics. In some cases, the estimates are so accurate that they are equal to the exact metric values. In this way the number of link metrics to be advertised is minimized without appreciably compromising the information contained in the link metrics.

The aggregation and distribution of link metrics for a subnetwork are typically executed by a representative node elected from the nodes in the subnetwork. This node is referred to as the subnetwork leader. The link metric advertisement is to be received by nodes outside the subnetwork. These nodes are referred to as external nodes. Upon receiving the advertisement, each external node is

responsible for decoding or estimating the values of the aggregated link metrics of the subnetwork.

FIG. 10, numeral 1000, is a flow chart of one embodiment of steps in accordance with the method of the present invention, which provides efficient aggregation of link metrics for a subnetwork in a communication network, comprising the steps of: A) selecting, by a predetermined subnetwork leader, a subset of the nodes in the subnetwork to be exposed nodes (1002); B) for each link metric, generating, by the subnetwork leader, a spanning tree virtual topology connecting exposed nodes (1004); and C) assigning link metric values to the virtual links on the spanning tree and advertising the metric values in a topology broadcast by the subnetwork leader (1006). Where selected, one may add, after step (A) a step of reducing, by the subnetwork leader, an original subnetwork topology, stored in a topology database of the subnetwork leader, to a full-mesh virtual topology wherein each virtual link connecting a pair of exposed nodes on the full-mesh virtual topology is derived utilizing a path selection optimization (1008). The advertised link metric values in step (D) typically include a number of predetermined specific link metric values for replacing selected estimates derived from each spanning tree.

Thus, the original topology of the communication subnetwork is reduced by the subnetwork leader to a full-mesh virtual topology, connecting a set of M predetermined exposed nodes, M a positive integer, such that the metric values on each virtual link on the full-mesh virtual topology are derived from a path selection optimization, e.g., a maximum bandwidth path or a minimum delay path. Then, for each link metric, the subnetwork leader

generates a spanning tree virtual topology from the full-mesh virtual topology for encoding the metric values on the full-mesh topology. The subnetwork leader then advertises the encoded link metric values, i.e., those
5 assigned to the virtual links on the spanning trees, to the external nodes via a topology broadcast. For each link metric, a full-mesh topology connecting the M exposed nodes is derived, by each external node, by means of a decoding technique, set forth below, from the spanning
10 tree topology such that the link metric value associated with each virtual link on the derived full-mesh topology is an estimate of the link metric value associated with the corresponding virtual link in the original full-mesh representation of the subnetwork topology.

15

Although generally the link metric values associated with a link are substantially symmetrical in both directions of a link, the method of the present invention is also applicable where link metric values are not
20 symmetrical. In the case of the generic link metrics bandwidth and delay, the link bandwidth may be taken to be the minimum bandwidth between the two directions, and the delay to be the maximum delay between the two directions. In the present invention, an estimate for
25 bandwidth is a lower bound, and where it is a lower bound on the minimum bandwidth between the two directions of a link, the bandwidth estimation is the lower bound on bandwidth for both directions. In the present invention, an estimate for delay is an upper bound, and where it is an
30 upper bound on the maximum delay between the two directions, the estimate for delay is the upper bound on delay for both directions.

The present invention provides for selectively
35 utilizing, for example, one of two types of path selection

optimization: bandwidth maximization and delay minimization. A path for the bandwidth maximization is determined utilizing a maximum bandwidth path algorithm. A path for the delay minimization is determined utilizing a
5 minimum delay path algorithm. The maximum bandwidth path algorithm is a straightforward variation of the shortest path algorithm. The present invention also utilizes a graph algorithm, the Maximum Spanning Tree algorithm, and its equivalent version Minimum Spanning
10 Tree. The Maximum/Minimum Spanning Tree algorithm is used to determine a spanning tree with the maximum/minimum sum of link weights.

The present invention combines the selected path
15 optimization and spanning algorithms in a unique manner to provide a new, efficient method for aggregation of link metrics. The path for the delay minimization is determined utilizing a Minimum Delay Link Metric Aggregation Heuristic. Where path optimization is based
20 on maximum bandwidth, the link metric aggregation algorithm in accordance with the present invention is referred to as the Maximum Bandwidth Link Metric Aggregation Heuristic. Where path optimization is based on minimum delay, the link metric aggregation algorithm
25 in accordance with the present invention is referred to as the Minimum Bandwidth Link Metric Aggregation Heuristic.

Thus, where link metrics are aggregated for bandwidth maximization utilizing the Maximum Bandwidth
30 Link Metric Aggregation Heuristic, the following steps, shown in the flow chart of FIG. 11, numeral 1100, of are utilized: A) determining a maximum bandwidth path for each pair of exposed nodes in the original subnetwork topology and representing each maximum bandwidth path

by a virtual link whose link metric values equal respective metric values of the maximum bandwidth path and wherein a collection of all virtual links constitutes the full-mesh virtual topology (1102); B) determining, from the full-mesh topology, a Maximum Bandwidth Spanning Tree using
5 a Maximum Spanning Tree algorithm and bandwidth as link weights and distributing the bandwidth of the virtual links in the resulting spanning tree to the external nodes in the network (1104); and C) determining, from the full-mesh
10 topology, a Maximum Delay Spanning Tree using a Maximum Spanning Tree algorithm and delay as link weights and distributing the delay of the virtual links in the resulting spanning tree to the external nodes in the network (1106).

15 The delay associated with a maximum bandwidth path connecting a pair of exposed nodes is an upper bound on the minimum delay for any path connecting the same pair of exposed nodes because a path requiring less bandwidth than the maximum bandwidth may always have
20 the maximum bandwidth path as an alternative.

The bandwidth of any virtual link not on the spanning tree is typically estimated by taking the minimum bandwidth among the virtual links along the unique path
25 that forms a loop with the virtual link, i.e., the bandwidth of the path. The maximum bandwidth of the virtual link not on the spanning tree cannot be larger than the estimate, or the tree will not be a maximum bandwidth spanning tree. On the other hand, it cannot be smaller than
30 the estimate, for there is a contradiction if there exists an alternate path with a larger bandwidth than the maximum bandwidth of the virtual link. Hence, the bandwidth of the unique path is indeed the same as the
35 bandwidth of the maximum bandwidth path or the corresponding virtual link in the full-mesh topology.

Where a Minimum Spanning Tree algorithm is utilized in the second step instead of a Maximum Spanning Tree algorithm, lower bound estimates are obtained by taking the maximum bandwidth among the virtual links along the
5 unique path.

The delay of any virtual link not on the spanning tree is typically estimated by taking the minimum delay among the virtual links along the unique path that forms a loop
10 with the virtual link, i.e., not the delay of the path since delay is an additive link metric. The delay of the virtual link not on the spanning tree cannot be larger than the estimate, or the tree is not a maximum delay spanning tree. Hence, the estimate serves as an upper bound on the
15 delay of the maximum bandwidth path or the corresponding virtual link in the full-mesh virtual topology.

Thus, as shown in FIG. 12, numeral 1200, where link metrics are aggregated for delay minimization utilizing
20 the Minimum Delay Link Metric Aggregation Heuristic, the following steps are included: A) determining a minimum delay path for each pair of exposed nodes in the original subnetwork topology wherein each minimum delay path is represented by a virtual link whose link metric values are
25 equal to the respective metric values of the minimum delay path and a collection of all the virtual links forms a full-mesh virtual topology (1202); B) determining, from the full-mesh topology, a Minimum Bandwidth Spanning Tree using a Minimum Spanning Tree algorithm and
30 bandwidth as link weights and distributing the bandwidth of the virtual links in the resulting spanning tree to the external nodes in the network (1204); and C) determining, from the full-mesh topology, a Maximum Delay Spanning Tree using a Maximum Spanning Tree algorithm and delay
35 as link weights and distributing the delay of the virtual

links in the resulting spanning tree to the external nodes in the network (1206).

5 The bandwidth associated with a minimum delay path connecting a pair of exposed nodes is a lower bound on the maximum bandwidth for any path connecting the same pair of exposed nodes because a path that can tolerate a larger delay than the minimum delay always has the minimum delay path as an alternative.

10

The bandwidth of any virtual link not on the spanning tree is typically estimated by taking the maximum bandwidth among the virtual links along the unique path that forms a loop with the virtual link, i.e., not the bandwidth of the path. The maximum bandwidth of the virtual link not on the spanning tree cannot be smaller than the estimate, or the tree would not be a minimum bandwidth spanning tree. Hence, the estimate serves as a lower bound on the bandwidth of the minimum delay path or the corresponding virtual link in the full-mesh topology.

20

The delay of any virtual link that is not on the spanning tree may be estimated by taking the minimum delay among the virtual links along the unique path that forms a loop with the virtual link, i.e., not the delay of the path since delay is an additive link metric. The delay of the virtual link not on the spanning tree cannot be larger than the estimate, or the tree is not a maximum delay spanning tree. Hence, the estimate serves as an upper bound on the delay of the minimum delay path or the corresponding virtual link in the full-mesh topology.

25

30

In both cases, the bounds on the link metrics of the virtual links may be somewhat flexible. When a particular bound on a link metric is flexible, one can additionally

35

advertise a specific virtual link to be used for replacing the estimate derived from the spanning tree. If additional metric values not associated with the spanning tree are advertised, it is important that they are indicated as such
5 in the advertisement, so that the spanning tree can be identified for generating metric estimates.

An example of the application of the invention using the Maximum Bandwidth Link Metric Aggregation Heuristic is shown in FIGs. 4-9, numerals 400, 500, 600, 700, 800
10 and 900 respectively. FIG. 4, numeral 400, is a representation of exemplary bandwidth and delay metrics (402) in a subnetwork for utilizing the method of the present invention. FIG. 5, numeral 500, is a representation
15 of exemplary bandwidth and delay metrics (502) for a full-mesh topology derived from maximum bandwidth paths utilizing the method of the present invention. FIG. 6, numeral 600, is a representation of exemplary metrics (602) for a maximum bandwidth spanning tree in
20 accordance with the method of the present invention. FIG. 7, numeral 700, is a representation of exemplary metrics (702) for a maximum delay spanning tree in accordance with the method of the present invention. FIG. 8, numeral 800, is a representation of exemplary metrics (802)
25 advertised for bandwidth and delay in accordance with the method of the present invention. FIG. 9, numeral 900, is a representation of exemplary estimates (902) for a full-mesh topology in accordance with the method of the present invention.

30

In the FIGs. 4-9, link metric values associated with a link are represented by an ordered pair, (x, y), where x is the bandwidth value and y is the delay value, and wherever x or y is replaced by an "", the corresponding link metric
35 value is not relevant in the context of the figure. Exposed

nodes A, B, C, and D are utilized to construct a full-mesh topology. In this example, the bandwidth values are accurately recovered as expected, and upper bounds on the delay values are obtained. Where the delay estimate for the virtual link AB differs considerably from the actual delay for the same virtual link, the delay value for that virtual link is typically additionally advertised together with an indication that it does not belong to the Maximum Delay Spanning Tree.

10

Thus, in one embodiment, shown in FIG. 13, numeral 1300, the method of the present invention for providing efficient aggregation of link metrics for a subnetwork in a communication network includes the steps of: A) selecting a subset of nodes in a subnetwork of the communication network (1302) to be exposed nodes; B) reducing the original subnetwork topology to a full-mesh virtual topology (1304) connecting the exposed nodes, such that each virtual link connecting a pair of exposed nodes on the full-mesh virtual topology is derived utilizing at least one of B1-B2: B1) determining a maximum bandwidth path in the original subnetwork topology and representing each maximum bandwidth path by a virtual link whose link metric values equal respective metric values of the maximum bandwidth path, and wherein a collection of all virtual links constitutes the full-mesh topology; B2) determining a minimum delay path in the original subnetwork topology and representing each minimum delay path by a virtual link whose link metric values equal respective metric values of the minimum delay path, and wherein a collection of all virtual links constitutes the full-mesh topology; C) where the full-mesh virtual topology is determined with the maximum bandwidth objective, C1) determining the spanning tree for encoding

and advertising the bandwidth metrics (1306) by means of:
C1a) determining, from the full-mesh topology, a Maximum
Bandwidth Spanning Tree using a Maximum Spanning Tree
algorithm and bandwidth as link weights; and C1b)
5 distributing the bandwidth of the virtual links in the
Maximum Bandwidth Spanning Tree to the external nodes in
the network; C2) determining the spanning tree for
encoding and advertising the delay metrics by means of:
C2a) determining, from the full-mesh topology, a Maximum
10 Delay Spanning Tree using a Maximum Spanning Tree
algorithm and delay as link weights; and C2b) distributing
the delay of the virtual links in the Maximum Delay
Spanning Tree to the external nodes in the network; D)
where the full-mesh virtual topology is determined with
15 the minimum delay objective, D1) determining the spanning
tree for encoding and advertising the bandwidth metrics
(1308) by means of: D1a) determining, from the full-mesh
topology, a Minimum Bandwidth Spanning Tree using a
Minimum Spanning Tree algorithm and bandwidth as link
20 weights; and D1b) distributing the bandwidth of the virtual
links in the Minimum Bandwidth Spanning Tree to the
external nodes in the network; D2) determining the
spanning tree for encoding and advertising the delay
metrics by means of: D2a) determining, from the full-
25 mesh topology, a Maximum Delay Spanning Tree using a
Maximum Spanning Tree algorithm and delay as link
weights; D2b) distributing the bandwidth of the virtual
links in the Maximum Delay Spanning Tree to the external
nodes in the network.

30

FIG. 14, numeral 1400, is a flow chart of one
embodiment of, in the step of assigning link metric values
to the virtual links in accordance with the present
invention, selecting a virtual link, on the full-mesh virtual

topology, whose metric value is to be estimated from the corresponding spanning tree, including the steps of: A)

5 determining whether the virtual link is on the spanning tree (1402); B) where the virtual link is on the spanning tree, a metric value of the virtual link on the spanning tree is taken to be an estimate for the metric value of the corresponding virtual link on the full-mesh topology (1404); C) where the virtual link is absent from the spanning tree, determining a unique path connecting
10 the exposed nodes attached at ends of the virtual link (1406), and determining whether the spanning tree is a maximum spanning tree (1408); C1) where the spanning tree is a maximum spanning tree, taking a minimum metric value among virtual links along a unique path to be an estimate for a metric value of the virtual link on the full-mesh virtual topology (1410); and C2) where the spanning tree is other than a maximum spanning tree, i.e., a minimum spanning tree, taking a maximum metric value among virtual links along a unique path to be an estimate
15 for a metric value of the virtual link on the full-mesh virtual topology (1412).
20

Compared to the prior art that divides nodes into master nodes and subordinate nodes such that the master
25 nodes maintain the network topology and carry out path computation on behalf of the subordinate nodes, the present invention provides the advantage of requiring little modification to existing routing protocols and architectures.

30

The present invention supports routing protocols that make use of link states, whereas in the prior art there are other methods that support routing protocols that make use of distance vectors. Aggregation of link states is

more tractable than aggregation of distance vectors since the latter requires inter-subnetwork cooperation.

5 The present invention requires advertisement of order M pieces of link state information when the number of exposed nodes is M . Unlike the approach that advertises a full-mesh representation of the subnetwork topology, the invention does not suffer from link metric explosion. The
10 estimated virtual topology proposed by the invention adequately reflects asymmetry in the original topology of the subnetwork.

15 While the prior art for routing to multiple destinations in a computer network tends to minimize the cost of communicating with multiple destinations, the present invention minimizes the number of link metrics to be advertised and estimates link metrics that are not advertised, thereby providing efficient routing.

20 The unique method of the present invention provides for representing link metric values on a full-mesh topology by a spanning tree topology and for estimating link metric values for links that are not part of the spanning tree, a capability that is particularly useful in
25 link state aggregation for asynchronous transfer mode network implementations in an internetworking environment.

30 I claim:

1. A method for providing efficient aggregation of link metrics for a subnetwork in a communication network, comprising the steps of:

5 A) selecting, by a predetermined subnetwork leader, a subset of nodes in the subnetwork to be exposed nodes,

B) for each link metric, generating, by the subnetwork leader, a spanning tree virtual topology connecting the exposed nodes, and

10 C) assigning link metric values to the virtual links on the spanning tree and advertising the metric values in a topology broadcast by the subnetwork leader.

2. The method of claim 1 inserting, between steps (1A) and (1B) a step of reducing, by the subnetwork leader, an original subnetwork topology, stored in a topology database of the subnetwork leader, to a full-mesh virtual topology wherein each virtual link connecting a pair of exposed nodes on the full-mesh virtual topology is derived utilizing a path selection optimization.

3. The method of claim 2 wherein at least one of 3A-3F:
10 3A) wherein the advertised link metric values include a number of predetermined specific link metric values for replacing selected estimates derived from each spanning tree;

3B) wherein a full-mesh virtual topology connecting
15 the exposed nodes is derived from the spanning tree topology such that a virtual link metric value associated with each virtual link on the full-mesh virtual topology is an estimate of a link metric value associated with a corresponding virtual link in the original full-mesh representation of the subnetwork topology;

20 3C) wherein the link metric values associated with a link are:

3C1) substantially symmetrical in both directions of a link, and

25 3C2) unsymmetrical in both directions of a link;

3D) wherein the link metric bandwidth is the minimum bandwidth between two directions and the link metric delay is the maximum delay between two
30 directions;

3E) wherein advertising the link metric values in a topology update includes advertisement of order M pieces of link state information when the number of exposed

nodes is M, where M is a positive integer that is less than or equal to the total number of nodes in the subnetwork;

3F) wherein assigning link metric values to the virtual links includes, for a given link metric, selecting a
5 virtual link, on the full-mesh virtual topology, whose metric value is to be estimated from the corresponding spanning tree utilizing the steps of 3F1-3F3:

3F1) determining whether the virtual link is on the spanning tree,

10 3F2) where the virtual link is on the spanning tree, a metric value of the virtual link on the spanning tree is taken to be an estimate for the metric value of the corresponding virtual link on the full-mesh topology,

3F3) where the virtual link is absent from the
15 spanning tree, determining a unique path connecting the exposed nodes attached at ends of the virtual link, and determining whether the spanning tree is a maximum spanning tree,

3F3a) where the spanning tree is a
20 maximum spanning tree, taking a minimum metric value among virtual links along a unique path to be an estimate for a metric value of the virtual link on the full-mesh virtual topology, and

3F3b) where the spanning tree is other
25 than a maximum spanning tree (i.e., a minimum spanning tree), taking a maximum metric value among virtual links along a unique path to be an estimate for a metric value of the virtual link on the full-mesh virtual topology.

30 4. The method of claim 2 wherein the path selection optimization is one of: bandwidth maximization and delay minimization.

5. The method of claim 4 wherein at least one of 5A-5B:

5A) wherein an estimate for a link metric bandwidth is a lower bound, and where the estimate is a lower bound on a minimum bandwidth between two directions of a link, the estimate for bandwidth is a lower bound on bandwidth
5 for both directions;

5B) wherein an estimate for a link metric delay is an upper bound, and where the estimate is an upper bound on the maximum delay between two directions, the estimate for delay is the upper bound on delay for both directions.
10

6. The method of claim 4 wherein:

6A) link metrics are aggregated for bandwidth maximization utilizing a Maximum Bandwidth Link Metric
15 Aggregation Heuristic, and link metrics are aggregated for delay minimization utilizing a Minimum Delay Heuristic, and

6B) a Maximum/Minimum Spanning Tree algorithm is utilized to determine a spanning tree with a
20 maximum/minimum sum of link weights.

7. The method of claim 6 wherein at least one of 7A-7C:

7A) wherein link metrics are aggregated for bandwidth maximization utilizing the Maximum Bandwidth
25 Link Metric Aggregation Heuristic that includes the steps of 7A1-7A3:

7A1) determining a maximum bandwidth path for each pair of exposed nodes in the original subnetwork topology and representing each maximum bandwidth path
30 by a virtual link whose link metric values equal respective link metric values of the maximum bandwidth path and wherein a collection of all virtual links constitutes the full-mesh virtual topology,

7A2) determining, from the full-mesh topology, a Minimum Bandwidth Spanning Tree using a Minimum Spanning Tree algorithm and bandwidth as link weights and distributing the bandwidth of the virtual links in the
5 resulting spanning tree to the external nodes in the network; and

7A3) determining, from the full-mesh topology, a Maximum Delay Spanning Tree using a Maximum Spanning Tree algorithm and delay as link weights and
10 distributing the delay of the virtual links in the resulting spanning tree to the external nodes in the network, and where further selected for 7A, including at least one of 7A4-7A5:

7A4) wherein the delay for bandwidth
15 maximization associated with the maximum bandwidth path connecting a pair of exposed nodes is an upper bound on the minimum delay for any path connecting the same pair of exposed nodes;

7A5) wherein the bandwidth of a non-spanning
20 tree virtual link is estimated by taking the minimum bandwidth among the virtual links along a unique path that forms a loop on the spanning tree with the virtual link, i.e., the bandwidth of the unique path;

7B) wherein, where a Minimum Spanning Tree
25 algorithm is utilized in step 6B, a lower bound estimate of the bandwidth of each non-spanning tree virtual link is obtained by determining the maximum bandwidth among the virtual links along a unique path that forms a loop on the spanning with the virtual link;

30 7C) wherein the delay of a non-spanning tree virtual link is estimated by determining a minimum delay among the virtual links along a unique path that forms a loop on the spanning tree with the virtual link.

8. The method of claim 6 wherein link metrics are aggregated for delay minimization utilizing the Minimum Delay Link Metric Aggregation Heuristic that includes the steps of 8A-8C:

5 8A) determining a minimum delay path for each pair of exposed nodes in the original subnetwork topology wherein each minimum delay path is represented by a virtual link whose link metric values are equal to the respective link metric values of the minimum delay path
10 and a collection of all the virtual links forms a full-mesh virtual topology,

 8B) determining, from the full-mesh topology, a Minimum Bandwidth Spanning Tree using a Minimum Spanning Tree algorithm and bandwidth as link
15 weights and advertising the bandwidths of the virtual links in the resulting spanning tree, and

 8C) determining, from the full-mesh topology, a Maximum Delay Spanning Tree using a Maximum Spanning Tree algorithm and delay as link weights and
20 advertising the delay of the virtual links in the resulting spanning tree;

and where further selected, at least one of 8D-8F:

 8D) wherein the bandwidth for delay minimization associated with a minimum delay path connecting a pair of
25 exposed nodes is a lower bound on the maximum bandwidth for any path connecting the same pair of exposed nodes;

 8E) wherein the bandwidth of a non-spanning tree virtual link is estimated by determining the maximum bandwidth along a unique path that forms a loop with the
30 virtual link, an estimate that serves as a lower bound on the bandwidth of the minimum delay path or a corresponding virtual link in the full-mesh virtual topology;

8F) wherein the delay of a non-spanning tree virtual link is estimated by determining the minimum delay along a unique path that forms a loop with the virtual link, an estimate that serves as an upper bound on the delay of the
5 minimum delay path or a corresponding virtual link in the full-mesh topology.

9. A method for providing efficient aggregation of link metrics for a subnetwork in a communication network having a plurality of interconnected subnetworks that consist of nodes and links, comprising the steps of:
- 5 9A) selecting a subset of the nodes in a subnetwork of the communication network to be exposed nodes,
- 9B) reducing the original subnetwork topology to a full-mesh virtual topology such that each virtual link connecting a pair of exposed nodes on the full-mesh virtual
- 10 topology is derived utilizing at least one of:
- 9B1) determining a maximum bandwidth path connecting each pair of exposed nodes in the original subnetwork topology and representing each maximum bandwidth path by a virtual link whose link metric values
- 15 equal respective metric values of the maximum bandwidth path and wherein a collection of all virtual links constitutes the full-mesh topology,
- 9B2) determining a minimum delay path connecting each pair of exposed nodes in the original
- 20 subnetwork topology and representing each minimum delay path by a virtual link whose link metric values equal respective metric values of the minimum delay path and wherein a collection of all virtual links constitutes the full-mesh topology;
- 25 9C) where the full-mesh virtual topology is determined with the maximum bandwidth objective, further utilizing the steps of:
- 9C1) determining the spanning tree for encoding and advertising the bandwidth metrics by means
- 30 of:
- 9C1a) determining, from the full-mesh topology, a Maximum Bandwidth Spanning Tree using a Maximum Spanning Tree algorithm and bandwidth as link weights, and

9C1b) distributing the bandwidth of the virtual links in the Maximum Bandwidth Spanning Tree to the external nodes in the network;

and

5 9C2) determining the spanning tree for encoding and advertising the delay metrics by means of:

9C2a) determining, from the full-mesh topology, a Maximum Delay Spanning Tree using a Maximum Spanning Tree algorithm and delay as link weights, and

10 9C2b) distributing the delay of the virtual links in the Maximum Delay Spanning Tree to the external nodes in the network;

9D) where the full-mesh virtual topology is determined with the minimum delay objective, further utilizing the steps of:

9D1) determining the spanning tree for encoding and advertising the bandwidth metrics by means of:

20 9D1a) determining, from the full-mesh topology, a Minimum Bandwidth Spanning Tree using a Minimum Spanning Tree algorithm and bandwidth as link weights, and

25 9D1b) distributing the bandwidths of the virtual links in the Minimum Bandwidth Spanning Tree to the external nodes in the network;

and

9D2) determining the spanning tree for encoding and advertising the delay metrics by means of:

30 9D2a) determining, from the full-mesh topology, a Maximum Delay Spanning Tree using a Maximum Spanning Tree algorithm and delay as link weights, and

9D2b) distributing the delay of the virtual links in the Maximum Delay Spanning Tree to the external nodes in the network.

10. The method of claim 9 wherein at least one of 10A-10B:

10A) wherein a Maximum/Minimum Spanning Tree algorithm is utilized to determine a spanning tree with a
5 maximum/minimum sum of link weights respectively; and

10B) wherein advertising the link metric values in a topology update includes advertisement of order M pieces of link state information when the number of exposed nodes is M, where M is a positive integer that is less than
10 or equal to the total number of nodes in the subnetwork.

1/6






FIG. 2

200



2/6




FIG. 4

400

FIG. 5

500





FIG. 6

600

3/6

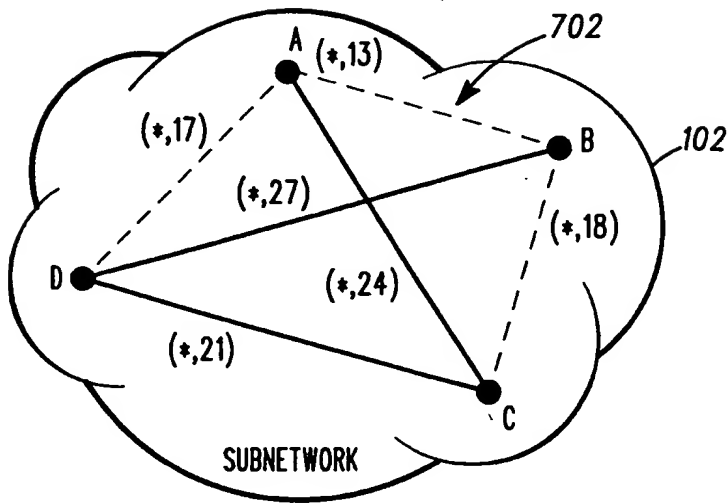


FIG. 7

700

FIG. 8

800



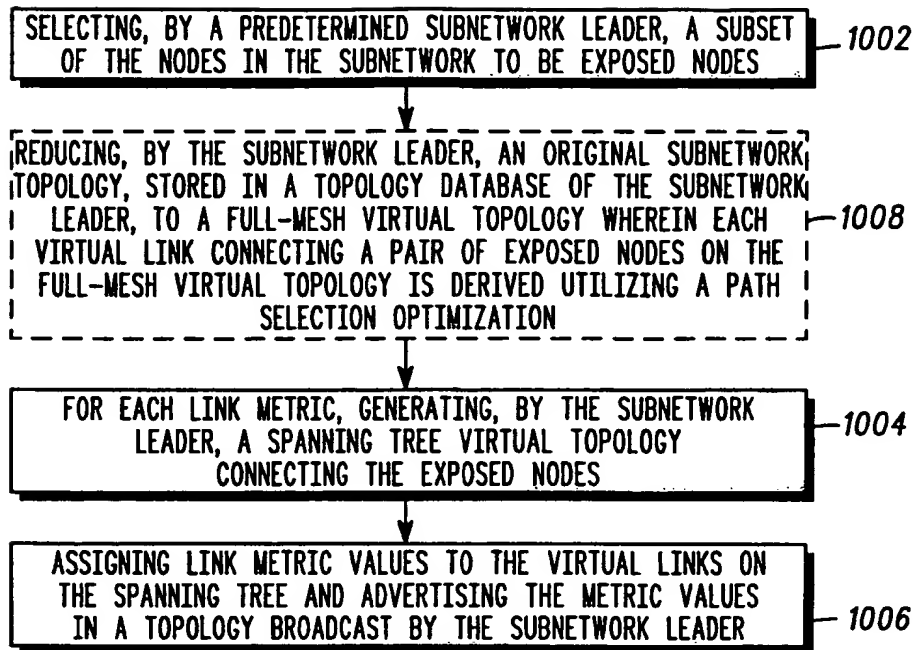
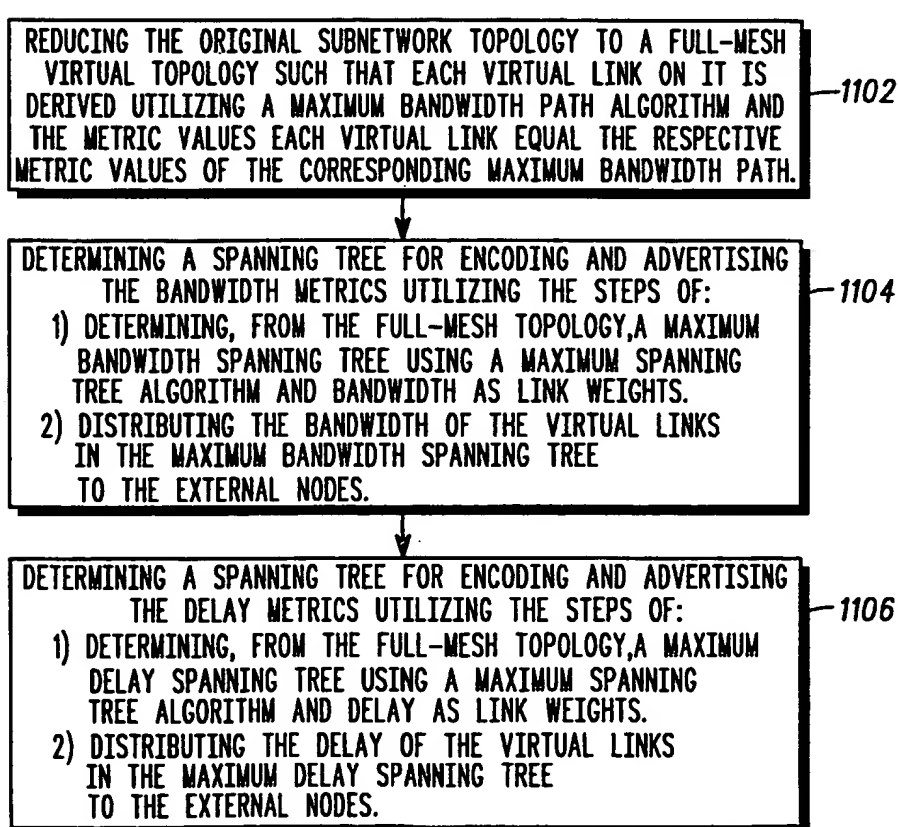


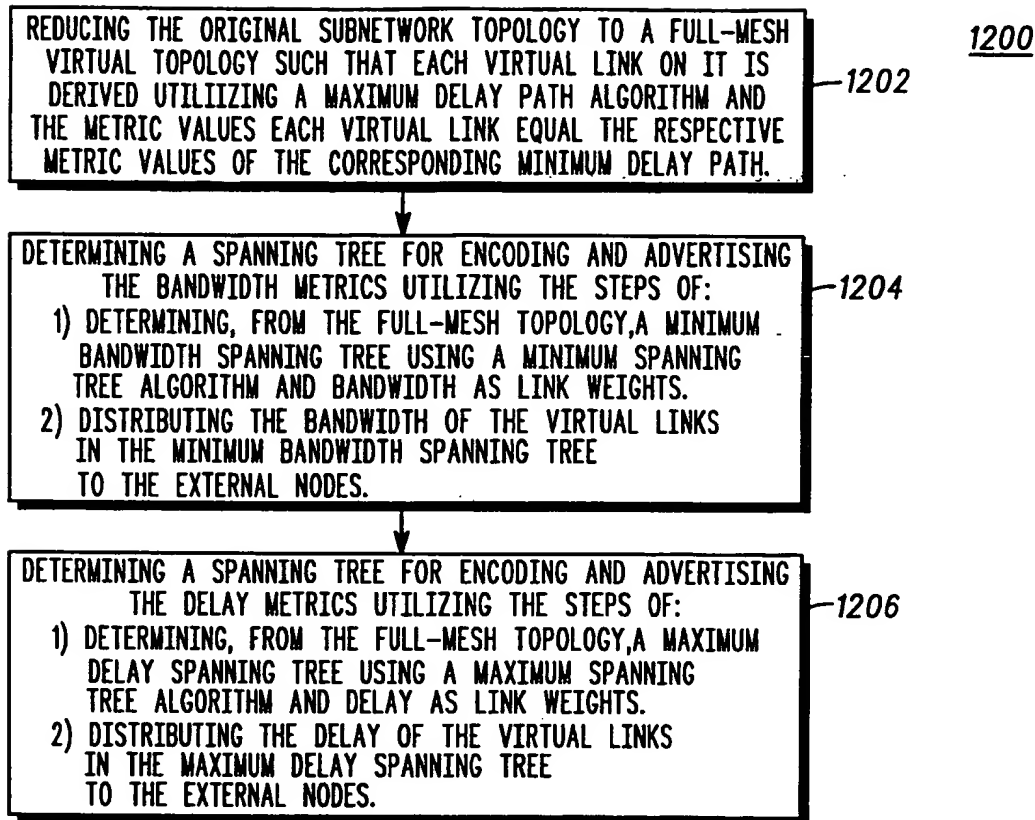
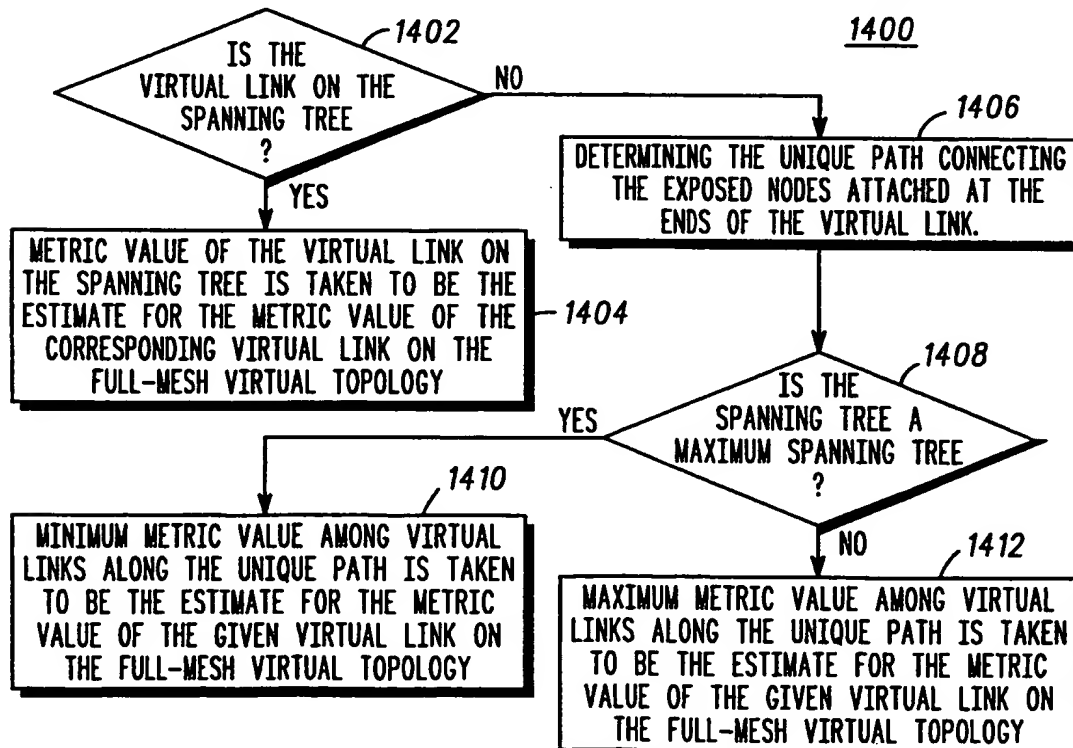
FIG. 9

900

4/6

FIG. 10**FIG. 11**

5/6

FIG. 12**FIG. 14**

6/6

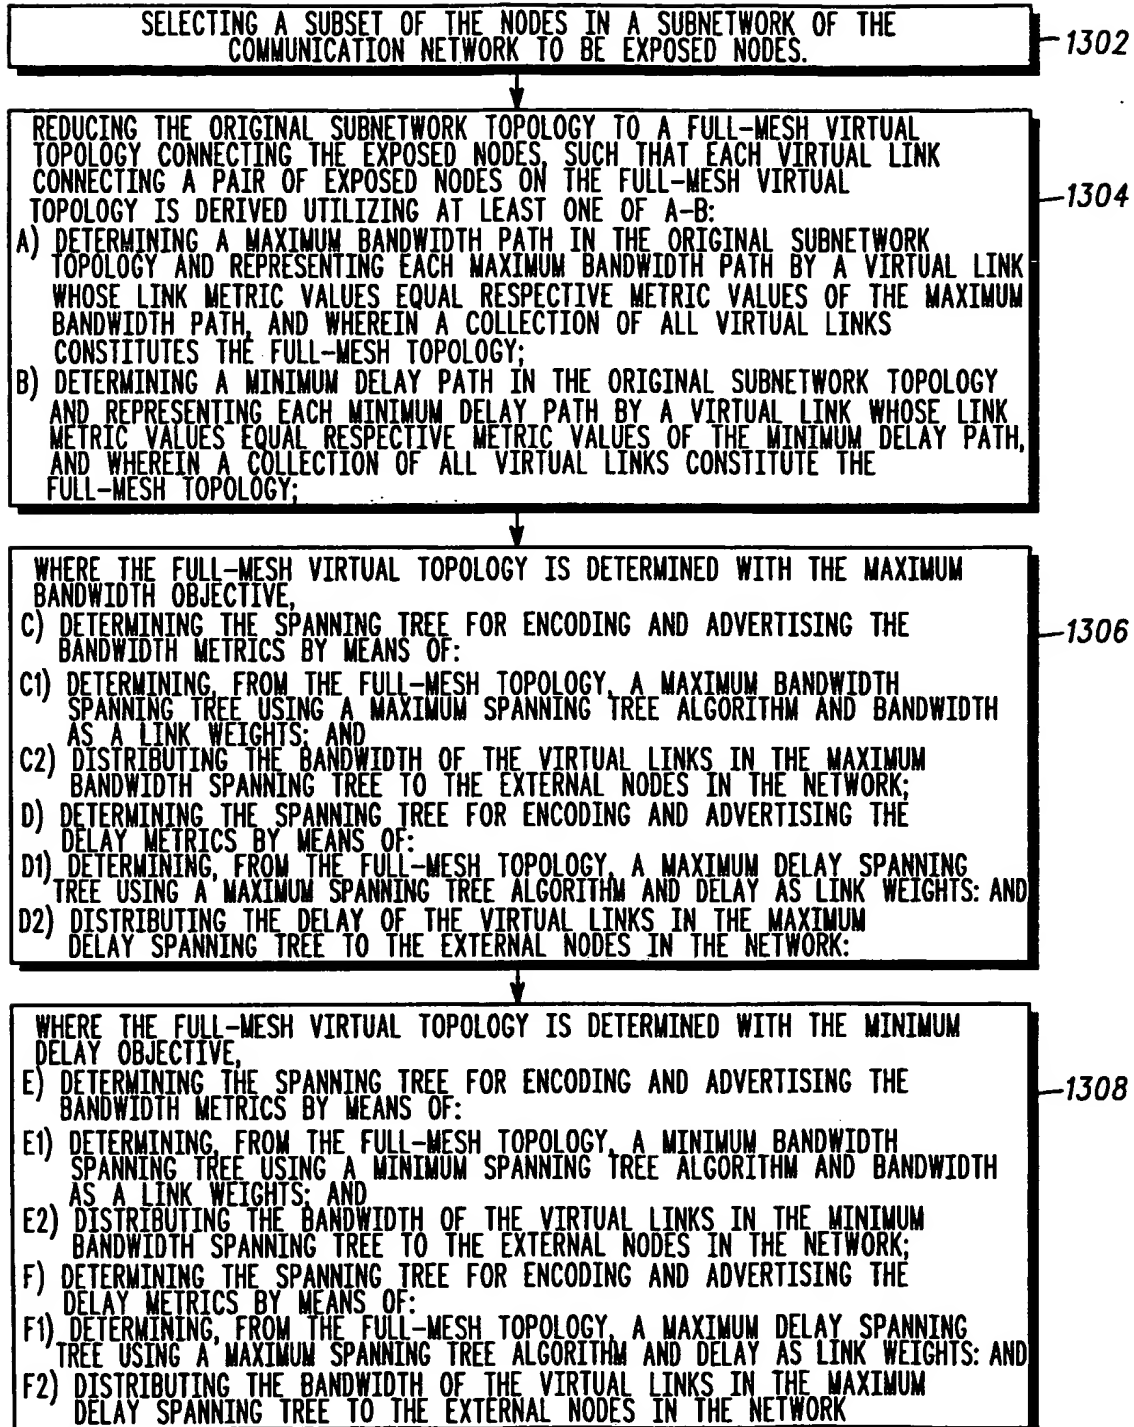


FIG. 13

1300

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US95/01789

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :H04L 12/44

US CL :370/54,94.3; 340/826

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/14, 17, 54, 58.2, 58.3, 60, 85.13, 94.3; 340/826, 825.03; 395/325

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US, A, 5,079,767 (Perlman) 07 January 1992, Figure 6, column 6, lines 32-44, column 11, lines 16-33.	1
Y	US, A, 5,234,592 (Perlman et al) 07 September 1993, Figure 2a, column 6, line 59- column 7, line 9.	1
A	US, A, 4,873,517 (Baratz et al) 10 October 1989, Figures 4 and 5, column 5, line 53 - column 6, line 34.	4, 6, 8
A	US, A, 5,042,027 (Takase et al) 20 August 1991, Figure 1, column 4, lines 41 - 54.	2

☐ Further documents are listed in the continuation of Box C.☐ See patent family annex.

* Special categories of cited documents:	T	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be part of particular relevance	X	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	Y	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	A	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

19 MAY 1995

Date of mailing of the international search report

31 JUL 1995

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

BENEDICT V. SAFOUREK

Telephone No. (703) 305-4364